

**IN THE UNITED STATES DISTRICT COURT  
FOR THE NORTHERN DISTRICT OF ILLINOIS  
EASTERN DIVISION**

KOVE IO, INC.

*Plaintiff,*

v.

GOOGLE LLC,

*Defendant.*

Case No. 1:23-cv-04244

Hon. Mary M. Rowland

---

**DEFENDANT GOOGLE LLC'S RULE 12(B)(6) MOTION TO DISMISS**

---

## TABLE OF CONTENTS

	<b>Page</b>
I. SUMMARY OF THE COMPLAINT.....	3
II. ARGUMENT.....	6
A. No Indirect Infringement; No Willfulness: Kove Does Not Allege Knowledge of the Asserted Patents During the Alleged Infringement Period. ....	6
B. The Complaint Does Not Plausibly Allege Direct Infringement by Google. ....	7
1. The '170 Patent (claim 15) and the '640 Patent (claim 10) require sending redirect messages, but the Complaint pleads a hierarchy- based approach.....	7
2. Claim 1 of the '170 Patent requires a hash function, but Kove cites exhibits showing the use of range-based identifiers, not hash-based identifiers. ....	10
3. Claim 17 of the '978 Patent requires a “transfer protocol,” but the Complaint does not identify one. ....	13
III. CONCLUSION.....	15

# **TABLE OF AUTHORITIES**

**Page(s)**

**Cases**

<i>Ashcroft v. Iqbal</i> , 556 U.S. 662 (2009).....	6, 7
<i>BASF Plant Sci., LP v. Commonwealth Sci. &amp; Indus. Rsch. Org.</i> , 28 F.4th 1247 (Fed. Cir. 2022) .....	6
<i>BOT M8 LLC v. Sony Corp. of Am.</i> , 4 F.4th 1342 (Fed. Cir. 2021) .....	7
<i>Kove IO, Inc. v. Amazon Web Servs., Inc.</i> , Dkt. 484 (N.D. Ill. Dec. 17, 2021) .....	1, 11
<i>Phillips v. Prudential Ins. Co. of Am.</i> , 714 F.3d 1017 (7th Cir. 2013) .....	9, 12
<i>SynQor, Inc. v. Artesyn Techs., Inc.</i> , 709 F.3d 1365 (Fed. Cir. 2013).....	6

**Other Authorities**

Federal Rule of Civil Procedure 12(b)(6) .....	6
--	---

Kove IO, Inc. (“Kove”) accuses Google of directly, indirectly, and willfully having infringed three related patents, each of which expired years ago. Kove alleges that two Google systems (called Colossus and Spanner) practiced four patent claims of the three Asserted Patents. The Court should dismiss Kove’s Complaint because Kove has not plausibly alleged that Google previously infringed any Asserted Claims. As discussed below, the Court can rely on Kove’s allegations to find no plausible infringement.

First, Kove’s claims for indirect infringement (*i.e.*, that Google allegedly induced or contributed to a third party’s infringement) and willful infringement must be dismissed because the Complaint acknowledges that Google never knew about any Asserted Patent before each expired. Because knowledge of the patents during the infringement period is a legal requirement of both indirect infringement and willfulness claims, those claims are not plausible in light of Kove’s factual allegations.

Second, Kove has not plausibly alleged direct infringement. For example, two asserted claims require a client to request the location of the data (*i.e.*, “location information”) from a “location server”;<sup>1</sup> if the location server does not have the location of the requested data, then the location server sends back a “redirect message” so that the client can ask the correct (different)

---

<sup>1</sup> The Court need not engage in any claim construction to grant this motion, but Google notes that certain terms of the Asserted Patents have been construed in a different case, and none of Google’s arguments in this motion are inconsistent with those constructions. In that case, this Court construed “client” as “a network-attached component (which may be software or hardware) that initiates update or lookup of identifier/location mappings from a location server with location request messages.” *Kove IO, Inc. v. Amazon Web Servs., Inc.*, Dkt. 484 at 21–23, 35 (N.D. Ill. Dec. 17, 2021). This Court construed “location information” in the ’978 Patent as “one or more identifiers and their associated locations” and in the ’170 and ’640 patents as “information pertaining to one or more locations of data and/or the identities of one or more location servers.” *Id.* at 12–19, 35. Lastly, this Court construed “location server” as “a network-attached component that maintains a set of identifier/location mappings that are modified or returned in response to location request messages from clients.” *Id.* at 19–21, 35.

location server. But, as confirmed by exhibits to the Complaint, Google’s accused products use a materially different hierarchy-based approach to retrieve the data instead of a “redirect” message. Accordingly, the facts pleaded cannot plausibly support the infringement allegations.

Another asserted claim requires location servers to organize and locate data based on a “hash function.” As explained more below, exhibits to the Complaint show that Kove accuses what is known as range-based identifiers, not hash-based identifiers. Both range-based and hash-based systems sub-divide large data tables and store the “splits” in separate locations, but they accomplish this in critically different ways. Consider a database storing data of U.S. presidents. In a range-based system, the presidents may be ordered chronologically and split by which century they first served in office. The “ranges” here would be 1789–1799, 1800–1899, etc. This range-based approach would result in four “splits” of unequal amounts. A hash-based system would take a different approach: an algorithm would equally distribute the presidents into each “split.” See below a visualization of this example:

	Number of Presidents in Each Split			
Splits	Split 1	Split 2	Split 3	Split 4
Range-Based	2	23	17	4
Hash-Based	12	12	11	11

The range-based approach prioritizes ease of retrieving information whereas the hash-based approach prioritizes reducing the “burden” on any given split by spreading the data more evenly. The Complaint’s exhibits clarify this important difference, dooming Kove’s infringement claims.

Finally, the remaining asserted claim requires that the Accused Products implement a claimed “transfer protocol.” Common examples of transfer protocols are FTP (file transfer protocol) and HTTP (hyper-text transfer protocol). Kove’s Complaint fails to include sufficient

factual allegations as to this required element, and instead provides only conclusory allegations parroting the claim language. Because this does not satisfy the requirement for fact-based pleading, this remaining claim should also be dismissed.

## **I. SUMMARY OF THE COMPLAINT<sup>2</sup>**

Kove accuses Google of having infringed US Patent Nos. 7,103,640 (the “’640 Patent”), 7,233,978 (the “’978 Patent”), and 7,814,170 (the “’170 Patent”) (Compl. ¶¶ 16–20, Dkt. 1) (collectively, the “Asserted Patents”). The Asserted Patents relate to storing and retrieving information in distributed networks. As discussed more below, the Asserted Patents disclosed alleged advancements by using “identification strings,” which specify the identity of “an entity,” and “location strings,” which specify the location of data associated with that entity. This information is stored in a distributed network using a hash function and retrieved by applying the hash function to an identifier string (*e.g.*, ’170 Patent, Claim 1). Or the information is retrieved by sending a requesting client the location string or a “redirect message” that tells the client where to find the location string (*e.g.*, ’170 Patent, claim 15; ’640 Patent, Claim 10). Alternatively, identifiers can be stored and transferred using a “transfer protocol” (*e.g.*, ’978 Patent, Claim 17).

The Complaint alleges direct and indirect infringement of claims 1 and 15 of the ’170 Patent, (Compl. ¶¶ 26–57); claim 17 of the ’978 Patent, (*id.* ¶¶ 58–77); and claim 10 of the ’640 Patent (*id.* ¶¶ 78–98) (collectively, the “Asserted Claims”). Kove accuses the Google products Spanner and Colossus (collectively, the “Accused Products”) of practicing the Asserted Claims. (*Id.* ¶ 22).<sup>3</sup>

---

<sup>2</sup> For this motion, Google accepts as true any well-pleaded factual allegations that are not contradicted by cited exhibits. To be clear, Google challenges the validity of the Asserted Claims.

<sup>3</sup> Kove alleges Spanner practices claim 1 of the ’170 Patent (Compl. ¶ 28); Colossus practices claim 15 of the ’170 Patent (*id.* ¶ 39); and Spanner and Colossus together practice claim 17 of the ’978 Patent (*id.* ¶¶ 61–77) and claim 10 of the ’640 Patent (*id.* ¶¶ 82–98).

Spanner is a relational database and Colossus is a cluster-level file system. (Compl., Ex. 88 at 2). At a very high level, a relational database like Spanner stores “structured data,” (*id.*, Ex. 10 at 1)—*e.g.*, data tables. Structured data may take the form of a table containing columns indicating a key identifying a particular file and additional columns listing the file’s metadata like “access permissions and data location.” (*Id.*, Ex. 8 at 3). Using components of the Colossus Control Plane, Colossus, in turn, provides clients with access to stored data, which includes storing documents and other file types in “D file servers.” (*Id.* at 3–4).

In part because the Spanner and Colossus datasets are massive, the datasets are split and stored at multiple locations. Spanner and Colossus are broken into regions, and further broken into multiple zones per region. (*id.*, Ex. 9). There are multiple ways to split datasets. Google’s “approach is to partition database tables into contiguous **key ranges** called splits.” (*Id.*, Ex. 15 at 1) (emphasis added); *see also id.*, Ex. 9; Ex. 17 at 5 (“Cloud Spanner distributes management of rows across multiple nodes by breaking up each table into several splits, **using ranges** of the lexicographically sorted primary key.”)). This redundancy allows Spanner and Colossus to function even if a zone were to fail. (*Id.*, Ex. 9).

As seen below, once the data is split, a split ID is assigned to each key range.

In this example, we have a table with a simple integer primary key.

Split	KeyRange
0	$[-\infty, 3)$
1	$[3, 224)$
2	$[224, 712)$
3	$[712, 717)$
4	$[717, 1265)$
5	$[1265, 1724)$
6	$[1724, 1997)$
7	$[1997, 2456)$
8	$[2456, \infty)$

(*Id.*, Ex. 15 at 3–4). Table splits can be referred to as “tablets.” (*Id.*, Ex. 19 at 4). Both user tables/tablets (*i.e.*, the table/tablets that store a user’s information) and user tablet location data can be split. (*Id.* at 4–5). As visualized below, user tablet location data is stored in metadata tablets. (*Id.* at 4). Sometimes, the queried metadata tablet does not know the user tablet’s location information or the location information is incorrect. In those instances, the client “recursively moves up the tablet location hierarchy” without any sort of prompt or message from the tablet. (*Id.* at 5). In other words, the client will continue to query other metadata tablets within the same metadata table set for the requested user tablet location. (*Id.*).

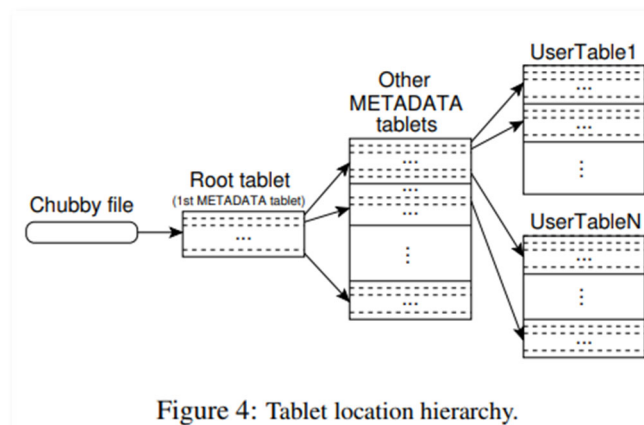


Figure 4: Tablet location hierarchy.



Using the above processes, the exhibits to the Complaint show Spanner and Colossus store data and data's location, split data into multiple locations using a range-based approach, and retrieve data for a user using a hierarchy approach to locate the data.

## II. ARGUMENT

Under Federal Rule of Civil Procedure 12(b)(6), a complaint must be dismissed if it fails to state a claim upon which relief may be granted. Fed. R. Civ. P. 12(b)(6). To survive a motion to dismiss, a complaint must contain sufficient factual matter to state a claim that is “plausible on its face.” *Ashcroft v. Iqbal*, 556 U.S. 662, 678 (2009) (citation omitted).

### A. No Indirect Infringement; No Willfulness: Kove Does Not Allege Knowledge of the Asserted Patents During the Alleged Infringement Period.

Kove fails to, and cannot, plausibly plead indirect infringement and willfulness because Kove alleges Google learned of the Asserted Patents only when it was served the Complaint, though the patents expired years earlier. Indirect infringement and willfulness allegations require knowledge of the patent at the time of infringement, so these claims must be dismissed.

“Liability for induced or contributory infringement under § 271(b) or (c) requires knowledge that the induced acts constitute patent infringement[,]” which “includes, in part, actual knowledge of the existence of the patent that is infringed.” *SynQor, Inc. v. Artesyn Techs., Inc.*, 709 F.3d 1365 (Fed. Cir. 2013) (quoting *Glob.-Tech Appliances, Inc. v. SEB S.A.*, 563 U.S. 754, (2011) (internal quotations omitted)). Similarly, a claim that a defendant willfully infringed a patent requires “that the accused infringer had a specific intent to infringe at the time of the challenged conduct.” *BASF Plant Sci., LP v. Commonwealth Sci. & Indus. Rsch. Org.*, 28 F.4th 1247, 1274 (Fed. Cir. 2022) (aff'g JMOL of no willfulness) (citation omitted).

Regarding Google's knowledge of the Asserted Patents, Kove asserts that “Google has had knowledge of the [’170/’978/’640] Patent since at least the time it was served with this Complaint.”

(Compl. ¶¶ 50, 73, and 94). Thus, Kove’s sole factual allegation is that Google did not learn of the Asserted Patents until 2023. But the Asserted Patents indisputably expired well before June 30, 2023. Accordingly, because Kove alleges only that Google knew of the Asserted Patents as of the time of the Complaint, Kove’s indirect infringement and willfulness claims must be dismissed.

**B. The Complaint Does Not Plausibly Allege Direct Infringement by Google.**

To meet the plausibility standard, Kove must plead “factual content that allows the court to draw the reasonable inference that the defendant is liable for the misconduct alleged.” *Iqbal*, 556 U.S. at 678. “[W]hile a patentee’s pleading obligations are not insurmountable, a patentee may subject its claims to early dismissal by pleading facts that are inconsistent with the requirements of its claims.” *BOT M8 LLC v. Sony Corp. of Am.*, 4 F.4th 1342, 1346 (Fed. Cir. 2021) (citation omitted).

**1. The ’170 Patent (claim 15) and the ’640 Patent (claim 10) require sending redirect messages, but the Complaint pleads a hierarchy-based approach.**

The allegations that Colossus infringed claim 15 of the ’170 Patent and Spanner and Colossus together infringed claim 1 of the ’640 Patent are implausible because the cited exhibits describe an altogether different method of finding information *without* any redirect message.

Claim 15 of the ’170 Patent requires “sending a redirect message;” Claim 10 of the ’640 Patent requires “transmitting a redirect message”:

Claim 15 of the ’170 Patent	Claim 10 of the ’640 Patent
A method of handling location queries in a network, the network comprising a plurality of location servers including data location information, the method comprising: correlating each one of a plurality of identifiers with at least one of a plurality of locations in the network, each one of the plurality of identifiers	A method for retrieving data location information for data stored in a distributed network, comprising the steps of: a) receiving at a first client a data query for retrieving data associated with an identification string, wherein the data is stored at a data repository in the distributed network and wherein a

<p>identifying a respective one of a plurality of data entities, wherein the data entities are stored in corresponding locations in the network;</p> <p>receiving a location query from a client at one of the plurality of location servers, the location query requesting location information identifying a location of a data entity included in the data entities;</p> <p>determining which of the plurality of location servers includes the location information;</p> <p>sending a location response message to the client in response to determining the one of the plurality of location servers includes the location information, the location response message comprising the location information; and</p> <p><b><i>sending a redirect message to the client in response to determining the one of the plurality of location servers fails to include the location information, the redirect message identifying which of the plurality of location servers includes the location information.</i></b></p> <p>(See '170 Patent, Claim 15) (emphasis added).</p>	<p>location string associated with the identification string of the data is stored in at least one of a plurality of data location servers;</p> <p>b) transmitting a data location request from the first client to a first data location server to retrieve the location string associated with the identification string in the data query, the data location request including the identification string;</p> <p><b><i>c) if the first data location server does not possess the location string, transmitting a redirect message to the first client, the redirect message containing information for use by the first client to calculate a location of a second data location server, wherein the second data location server contains the location string;</i></b></p> <p>d) calculating the location of the second data location server at the first client; and</p> <p>e) transmitting the data query from the first client to the second data location server.</p> <p>(See '640 Patent, Claim 10) (emphasis added).</p>
--	---

As shown above, both asserted claims require that if a queried location server or data location server does not have the required location information, then ***that server*** transmits or sends a “redirect message to [the/the first] client.” Furthermore, the claims require that the redirect message (1) identifies a location server that includes the location information ('170 Patent) or (2) contains information the client can use to calculate the location of another server that contains the location information ('640 Patent). Yet Kove does not allege that the Accused Products send or transmit a “redirect message” to a client, let alone a redirect message containing the specifically claimed information. To the contrary, Kove cites the following screenshot of Exhibit 19, which

demonstrates that the requesting clients (as opposed to the responding servers) of the Accused Products run a “location algorithm” to “recursively move up the tablet location hierarchy” to find the location:

The client library caches tablet locations. If the client does not know the location of a tablet, or if it discovers that cached location information is incorrect, then it recursively moves up the tablet location hierarchy. If the client's cache is empty, the location algorithm requires three network round-trips, including one read from Chubby. If the client's cache is stale, the location algorithm could take up to six round-trips, because stale cache entries are only discovered upon misses (assuming that METADATA tablets do not move very frequently). Although tablet locations are stored in memory, so no GFS accesses are required, we further reduce this cost in the common case by having the client library prefetch tablet locations: it reads the metadata for more than one tablet whenever it reads the METADATA table.

**Source:** Compl., Ex. 19 at 5.

(Compl. ¶¶ 48, 90, and 91) (emphasis added). Although Kove alleges that “the location algorithm instructs the client to redirect and access another tablet server to find the location of the tablet,” (Compl. ¶ 48),<sup>4</sup> that allegation contradicts Exhibit 19, which Kove is relying on, so the Court does not need to credit the allegation. *Phillips v. Prudential Ins. Co. of Am.*, 714 F.3d 1017, 1020 (7th Cir. 2013) (“To the extent that an exhibit attached to or referenced by the complaint contradicts the complaint's allegations, the exhibit takes precedence.”) (citation omitted).

More specifically, Exhibit 19 discloses that a client “recursively moves up the tablet location hierarchy,” which means the “location algorithm” controls where the client looks for the information according to a hierarchy, not a redirect message from a server. Using the claimed

---

<sup>4</sup> Similarly, for the '640 Patent, Kove alleges that the “location algorithm redirects to another tablet location, utilizing a location algorithm which calculates a location of second data location and the client (i.e. Spanner and Colossus) makes this calculation.” (Compl. ¶¶ 90 and 91).

“redirect message” would not only contradict the hierarchy approach described in Exhibit 19, but it would be inconsistent with Exhibit 19’s reference to “six round-trips.” If the Accused Products were using the claimed “redirect message,” it would point the client to the correct location, making “six round-trips” unnecessary.

Kove’s allegations, which are contradicted by the attached exhibits, thus cannot plausibly support a claim of infringement. In short, the claims require a server sending a redirect message to a client, and that the redirect message provides the information necessary to identify the correct data location. Exhibit 19 to the Complaint demonstrates that is not how the Accused Products work. Instead, the clients use a hierarchy-based approach. These are two entirely distinct methods of identifying a data location, such that Kove’s claims can and should go no further.

**2. Claim 1 of the ’170 Patent requires a hash function, but Kove cites exhibits showing the use of range-based identifiers, not hash-based identifiers.**

The allegation that Spanner infringes claim 1 of the ’170 Patent is not plausible because Kove alleges infringement by use of hash-based identifiers whereas the cited exhibits explicitly show the use of range-based identifiers.

Claim 1 of the ’170 Patent recites using a “hash function” to organize and retrieve data:

A system for managing data stored in a distributed network, the system comprising:  
a data repository configured to store a data entity, wherein an identifier string identifies the data entity; and  
a data location server network comprising a plurality of data location servers, wherein data location information for a plurality of data entities is stored in the data location server network, at least one of the plurality of data location servers includes location information associated with the identifier string, each one of the plurality of data location servers comprises a processor and a portion of the data location information, *the portion of the data location information included in a corresponding one of the data location servers is based on a hash function used to organize the data location information across the plurality of data location servers*, and each one of the data location servers is configured to determine the at

least one of the plurality of data location servers based on *the hash function applied to the identifier string*.<sup>5</sup>

(See '170 Patent, Claim 1) (emphasis added). Kove does not plausibly allege infringement because it cites exhibits showing Spanner using range-based identifiers (“key ranges”), not hash-based identifiers—a fact that even Kove admits. (See Compl. ¶ 32 (“Google Spanner divides user data, or data entities, into chunks called splits, and these splits are stored at different Cloud Spanner data repository servers present at different locations *based on key ranges*.”) (emphasis added); *id.* (Exhibit 15 Figure showing a table split using key ranges); Compl. ¶ 36 (“Spanner partitions tables into contiguous key ranges called splits and divide [sic] data among servers by key ranges.”)).

Outside the exhibits, Kove’s allegations just parrot the claim language and note that hash-based identifiers are referenced within Exhibits 15-18. (Compl. ¶¶ 33, 34, 36–38). But the exhibits reference hash-based identifiers only as one among various techniques that *users* can implement to avoid hotspots attributable to the fact that Spanner divides data by key ranges. (Compl. ¶ 33 and 36 (both citing Ex. 16, which explains that “Cloud Spanner divides data among servers by key ranges, which means your inserts will be directed at a single server, creating a hotspot.”)).<sup>6</sup> This demonstrates not only that the accused Spanner product used range-based identifiers, but also that

---

<sup>5</sup> Like before, *see supra* note 1, claim construction is not necessary and Google’s motion is consistent with constructions this Court made in a different case involving the Asserted Patents. In that case, this Court construed the phrase “based on a hash function used to organize the data location information across the plurality of data location servers . . . based on the hash function applied to the identifier string” as “the portion of the data location information included in a corresponding one of the data location servers is based on a hash function that maps identifier strings to one or more of the data location servers, and each one of the data location servers is configured to determine the at least one of the plurality of data location servers based on the hash function applied to the identifier string.” *Kove v. Amazon Web Servs.*, Dkt. 484 at 24–30, 35.

<sup>6</sup> Key ranges are mentioned again in Exhibit 16—“A split holds a range of contiguous rows”—when describing how Spanner “splits” data (Compl. Ex. 16 at 3; *id.* ¶ 34). Exhibit 16 is replete with the examples that primary keys are sorted “contiguously” (i.e. next in a sequence), which is a defining characteristic of range-based identifiers. (*Id.*, Ex. 16 at 1, 7, and 9). Exhibit 16 therefore demonstrates that Spanner used “key ranges,” and not hashes.

this design might lead users to implement techniques such as hashing keys to avoid hotspots. In any event, the fact that Google’s documentation to users identified hashing keys as one technique to avoid hotspots is beside the point because allegations that Google induced or contributed to Google users’ alleged infringement is an allegation of indirect infringement, not direct infringement, and must be dismissed for the reasons stated in Section II.A. Again, the exhibits control, *Phillips v. Prudential*, 714 F.3d at 1020, and they show Spanner using range-based identifiers, not hash-based identifiers.

Kove’s citation to Exhibit 17 similarly shows no plausible infringement. Exhibit 17 teaches a limited context wherein *users can utilize hashing* when they decide to implement an “integer sequence as a key” (also relevant to avoiding hotspots caused by Spanner’s range-based approach), and goes so far as to say that this approach is “not recommended.” (Compl., Ex. 17 at 8–10). Again, Exhibit 17 makes clear that Spanner used ranges, not hashes. For example, under the header “Table splits and key choice,” Exhibit 17 states that “Cloud Spanner distributes management of rows across multiple nodes by breaking up each table into several splits, *using ranges of the lexicographically sorted primary key.*” (*Id.*, Ex. 17 at 3–4) (emphasis added). Further, Exhibit 17 uses the below figure as an example, which depicts the use of “Key Range[s].”

Key Range (inclusive)	Row Count in Split	Split ID	Split Leader Node (ID, Zone)
$-\infty$ to 100	100 (rows 1-100)	0	1 a
101 to 200	100	1	2 a
201 to 300	100	2	1 b
301 to 400	100	3	2 b
401 to 500	100	4	1 c
501 to $\infty$	100 (501-600)	5	2 c

(*Id.* at 5–6).



Kove cites Exhibit 15 for the allegation that after a table is split, the data is “assigned to different Spanner servers based on key ranges, such as hashed\_id and id.” (Compl. ¶ 36). Exhibit 15 outlines the salient differences between range- and hash-based approaches, explaining the range-based approach is the one used by Spanner, but that users can write their applications to instead implement a hash-based approach:

Another example is [Spanner’s] data layout, where *we use range sharding*. Applications that do range scans can use this layout to get high performance. But some applications don’t need range scans; and *if those applications don’t want to worry about the performance cost of potential row-range hotspots, they can put a hash value in their primary key, effectively causing Spanner to use hash sharding*. In general, Spanner has been designed to offer powerful tools to application builders, but to *give those builders a high degree of control* over any tradeoffs between powerful functionality and performance.

(*Id.*, Ex. 15 at 10) (emphasis added).

All Exhibits show the same thing: Spanner using range-based identifiers. Thus, once allegations that contradict the Complaint’s exhibits are set aside, all that is left is an allegation that Google provided a system that used range-based identifiers. At most, Kove alleges that users can alter that system. But Kove has pleaded no viable direct infringement claim against Google.

### **3. Claim 17 of the ’978 Patent requires a “transfer protocol,” but the Complaint does not identify one.**

The allegation that Spanner and Colossus together infringed claim 17 of the ’978 Patent is not plausible because Kove does not allege a “transfer protocol” used by Spanner or Colossus.

Claim 17 of the ’978 Patent recites providing a “transfer protocol” to transport identifier and location information, and storing the location information in the format of the protocol:

A method of scaling at least one of capacity and transaction rate capability in a location server in a system having a plurality of location servers for storing and retrieving location information, wherein each of the plurality of location servers stores unique set of location information of an aggregate set of location information, the method comprising:



*providing a transfer protocol configured to transport identifier and location information, the location information specifying the location of information related to the identifier;*

*storing location information formatted according to the transfer protocol at a first location server;*

receiving an identifier and a location relevant to the identifier at the first location server;

storing the received location in a location store at the first data location server, the location store comprising a plurality of identifiers, each identifier associated with at least one location, wherein the received location is associated with the received identifier in the location store; and

transferring a portion of the identifiers and associated locations to a second data location server when a performance criterion of the first location server reaches a predetermined performance limit.

(See '978 Patent, Claim 17) (emphasis added). Well-known transfer protocols like FTP (file transfer protocol) and HTTP (hyper-text transfer protocol) describe not just the fact that files and hyper-text are transferred, but describe a protocol for how transfer happens. Here, although the Complaint meanders through various Google systems and concepts—Spanner, Colossus, BigTable, Rowkeys, SSTables, Metadata tables, Chubby, and tablets and tablet servers—it does not allege which provides the transfer protocol or stores location information according to the transfer protocol. (Compl. ¶¶ 62–67).

At most, Kove alleges that information is transferred; however, even then, Kove is inconsistent on what systems and information allegedly meet the claimed method steps relating to a transfer protocol. First, Kove makes the bare assertion that the “Accused Products provide a transfer protocol configured to transport identifier and location information.” (*Id.* ¶ 62). Incongruously, Kove supports this claim by quoting Exhibit 18, which relates to how Spanner and Colossus store information using tablets. (*Id.*). There are no facts pleaded about any transfer protocol. Next, Kove alleges that Colossus alone provides the transfer protocol “on a request to recover a tablet.” (*Id.* ¶ 63). Here, too, Kove points to how data is stored rather than anything about

how data is transferred or what provides the transfer protocol. (*Id.* (citing Exs. 8 and 19)). Kove then seemingly changes horses in midstream to allege that “retriev[ing] data from [a] table” is the “exemplary transfer protocol” and that this is completed by “BigTable us[ing] Rowkeys.” (Compl. ¶ 67).

These random, inconsistent allegations do not add up to a coherent claim, much less put Google on notice of what is accused of infringement. Recognizing that Kove need not prove its claim at this stage of the case, more is required as a matter of pleading under *Iqbal* and *Twombly*. When the asserted claim requires “providing a transfer protocol configured to transport identifier and location information” and “storing location information formatted according to the transfer protocol,” it does not suffice as a matter of pleading to just say the Accused Products store and retrieve data. Kove must plead facts sufficient to support its allegations about a transfer protocol.

### **III. CONCLUSION**

For the foregoing reasons, Google respectfully requests that the Court dismiss Kove’s Complaint for failure to state a claim upon which relief can be granted. Further, Kove should be given leave to replead its allegations regarding claim 17 of the ’978 Patent only if it makes specific and plausible allegations related to the Accused Products addressing the claimed “transfer protocol.”

Dated: September 26, 2023

Respectfully submitted,

/s/ Robert W. Unikel

Robert W. Unikel (Bar No. 6216974)

Douglas L. Sawyer (Bar No. 6275849)

Mark T. Smith (Bar No. 6315040)

John A. Cotiguala (Bar No. 6311056)

Grayson S. Cornwell (Bar No. 6338845)

Summer C. Stevens (Bar No. 6342946)

**PAUL HASTINGS LLP**

71 S. Wacker Drive, Suite 4500

Chicago, Illinois 60606

Tel: (312) 499-6000

Fax: (312) 499-6100

robertunikel@paulhastings.com

dougsawyer@paulhastings.com

marksmith@paulhastings.com

johncotiguala@paulhastings.com

graysoncornwell@paulhastings.com

summerstevens@paulhastings.com

*Counsel for Defendant,*

GOOGLE LLC

**CERTIFICATE OF SERVICE**

I hereby certify that a true and correct copy of the foregoing document was filed electronically in compliance with Fed. R. Civ. P. 5(a) on September 26, 2023. As of this date, all counsel of record had consented to electronic service and are being served with a copy of this document through the Court's CM/ECF system under Fed. R. Civ. P. 5(b) and (c).

/s/ Robert W. Unikel

Robert W. Unikel